AI GOVERNANCE SERIES | REGULATORY PERSPECTIVE

Colorado Built a Landmark AI Law, Then Dismantled It Before It Took Effect.

The lesson is not about Colorado. It is that AI governance must survive reversals, not just additions.

By Greg Aldrich | Global CIO & Strategic Advisor | June 2026

I live and work in Colorado, so the state's AI law is not an abstract regulatory development to me. It is a local signal of a much larger governance shift, and the story it tells is more instructive than the one I expected to write.

My previous article framed AI governance as a multi-state operating model problem, using Connecticut's omnibus law to show how requirements are fragmenting across jurisdictions. Colorado was supposed to be the next clean example in that pattern: another comprehensive state law, another set of obligations to map. But Colorado did not follow the pattern. It inverted it.

Colorado built the first comprehensive AI law in the United States, then repealed and replaced it before it ever took effect. That is a different kind of warning than Connecticut, and arguably a more important one. Connecticut showed that the regulatory surface is expanding. Colorado shows that it is also volatile, and that governance programs built to a specific law can be invalidated not only by new requirements, but by the reversal of the requirements they were built for.

> ***"Connecticut showed that AI regulation is fragmenting. Colorado shows that it is volatile. Governance must survive reversals, not just additions."***

What Actually Happened in Colorado

The sequence matters, because the lesson is in the sequence rather than in any single provision.

Date	What Happened
May 2024	SB 24-205 signed into law. First comprehensive US state AI law, modeled on the EU AI Act. Original effective date: February 1, 2026.
Aug 2025	After failed amendment negotiations in a special session, SB 25B-004 delays the effective date to June 30, 2026.
Dec 2025 – Jan 2026	Federal executive order and a new DOJ AI Litigation Task Force signal federal pressure against state AI laws.
Mar 2026	Governor's AI Policy Working Group releases a draft framework to repeal and replace the law.
Apr 2026	A federal magistrate judge stays enforcement after the Attorney General stipulates, citing the pending rewrite.
May 14, 2026	SB 26-189 signed. Repeals and replaces SB 24-205 entirely. New effective date: January 1, 2027.
Result	The original Colorado AI Act never took effect. Organizations that prepared for it spent two years building toward a law that was replaced before it became operational.

Colorado's original AI Act, SB 24-205, was signed in May 2024 and was genuinely ambitious. It imported the conceptual architecture of the EU AI Act: a risk-based regime built on duties of reasonable care, risk management programs, annual impact assessments, and high-risk system classification, all aimed at protecting consumers from algorithmic discrimination in consequential decisions. It was widely regarded as the most comprehensive state AI law in the country.

It never became operational. Governor Polis signed it in 2024 with explicit reservations, openly inviting the legislature to revise it. Industry opposition was intense. A special session in August 2025 failed to reach a compromise, simply delaying the effective date. A federal executive order and a new Department of Justice AI Litigation Task Force added pressure against state AI laws. A social media and AI company sued, arguing the law was unconstitutional. In April 2026, a federal magistrate judge stayed enforcement. And in May 2026, SB 26-189 repealed and replaced the original framework entirely.

What Replaced It

The replacement is not a tweaked version of the original one. It is a different regulatory philosophy.

SB 26-189, effective January 1, 2027, abandons the EU-style risk-management model and adopts a disclosure-and-rights framework focused on automated decision-making technology. The high-risk AI classification is gone. The mandatory risk management programs are gone. The annual impact assessments are gone. The duty of care is gone. In their place are obligations built around consumer notice when automated technology materially influences a consequential decision, adverse-outcome explanations within 30 days, meaningful human review, and developer documentation.

This moves Colorado away from the EU model and toward the disclosure-and-consumer-rights approach that California has anchored. The two dominant US state approaches are now converging on notice and rights rather than risk management and duties of care, while the EU AI Act, even with its own high-risk obligations now delayed to 2027, is still a substantively different risk-management regime.

For an organization that spent 2024 and 2025 building toward Colorado's original framework, this is not a minor adjustment. The impact assessment infrastructure, the high-risk classification work, and the duty-of-care risk management programs were all built for a law that no longer exists. Some of that work is transferable. Much of it has to be reconceived for a fundamentally different compliance model.

> ***"An organization that built to Colorado's original AI Act spent two years preparing for a law that was repealed before it took effect. That is governance strategy drift in its most expensive form."***

This Is Not a Legal Problem First

The instinct, when a law like this changes, is to treat it as a legal problem: what does Legal now need us to do? That is the wrong first question.

The better question is architectural: can our operating model identify, classify, monitor, document, and defend AI-enabled decisions across every jurisdiction in which we operate, and adapt as those jurisdictions change direction?

This distinction is the whole argument of this series. Compliance is downstream of architecture. An organization with a mature AI governance operating model, an accurate system inventory, clear ownership, documented decisions, jurisdictional mapping, and experience of the Colorado reversal as a readjustment. An organization that had built a Colorado-specific compliance program experienced it as a write-off.

The Colorado story is the clearest illustration I have seen of why governance must be built as adaptable infrastructure rather than as a response to a specific law. The organizations that handled this well were not the ones with the best Colorado compliance program. They were the ones whose governance architecture was designed to absorb regulatory change without being rebuilt each time the regulations moved.

Colorado in the Wider Patchwork

Colorado is one data point in a landscape that is fragmenting and shifting simultaneously. Without repeating the detail from my Connecticut piece, the broader picture now includes:

- Connecticut SB 5, the omnibus approach covering AI companions, frontier models, employment decision tools, provenance, and social media, with staggered effective dates through 2028.
- Utah SB 149 focused on consumer protection and disclosure of generative AI interactions.

- Texas TRAIGA, denoting a broad AI governance framework with prohibited uses and state oversight.
- California AB 2013 and SB 942 focused on training data transparency plus AI-generated content transparency, anchoring the disclosure-and-rights model that Colorado has now joined.
- New York City Local Law 144 is a local regulation requiring bias audits and candidate notices for automated employment decision tools.
- Illinois's AI Video Interview Act requires disclosure, consent, and deletion rights for AI-analyzed video interviews.

The pattern across these is not convergence toward a single standard. It is divergence in multiple directions at once, with individual states also reversing course mid-stream. An organization operating nationally is not tracking a stable set of requirements that expands over time. It is tracking a moving, occasionally self-contradicting target.

The Federal Layer Does Not Resolve This

It would be reasonable to hope that federal action would harmonize this splintering. So far, it has not, and the federal posture has arguably amplified volatility rather than reducing it.

The federal picture remains guidance-heavy. The NIST AI Risk Management Framework is a voluntary framework for trustworthy AI, useful for structure but not binding. Federal AI guidance for agencies has continued to evolve, with the current administration's OMB guidance underscoring innovation, governance, and public assurance. But guidance for federal agencies does not eliminate state-by-state compliance exposure for private organizations.

More significantly, the federal government has begun actively challenging state AI laws. The DOJ AI Litigation Task Force intervened in litigation against Colorado's original law, and federal recommendations have urged Congress toward a unified, preemption-oriented framework. Whether or not federal preemption eventually materializes, the current effect is to add another layer of uncertainty: organizations now have to consider not only what each state requires, but whether those state requirements will survive federal challenge.

This is the opposite of the stability that would justify building a governance program around any single regulatory framework, state or federal.

The Global Overlay

For multinational organizations, the picture is even more complex. The major international frameworks are themselves diverging:

- The EU AI Act is still a substantive risk-management regime built on risk tiers, even though its high-risk obligations are now delayed until 2027.
- The UK has taken a context-based, pro-innovation posture rather than a single comprehensive AI statute.

- China has had binding generative AI measures for public-facing services in effect since 2023.
- Canada's AIDA, proposed under Bill C-27, remains unsettled.
- Brazil is advancing AI regulation through its own legislative process.

The practical consequence is that a multinational organization cannot treat AI governance as a US compliance project with an international addendum. AI governance is becoming a global regulatory translation problem, in which the same AI system must be governed by materially different and independently shifting frameworks in every market where it operates.

The Innovation Tension, Honestly Stated

There is a genuine tension here that deserves to be named without taking a political side. The Common Sense Institute, a Colorado think tank, argued that the state's original AI policy could impose high economic costs, including projected job losses and reduced output. I cite this as a viewpoint in the debate, not as a settled fact, and the organization has a clear policy orientation.

But the core tension is real and worth engaging with directly. The question is not regulation versus innovation, which is a false binary. The real question is whether regulation can be operationalized without creating enough friction to slow responsible adoption. Colorado's reversal was driven in significant part by the perception that the first law's compliance burden was disproportionate to its benefits. Whether that perception was accurate or not, it was decisive.

For enterprise leaders, the lesson is that regulations that are difficult to operationalize tend to be politically unstable. Building a governance program around a heavy, contested compliance regime carries the risk that the regime itself will not survive, exactly as happened in Colorado.

What Good Looks Like

The governance capabilities that allow an organization to absorb regulatory volatility rather than being whipsawed by it are consistent regardless of which way any individual law moves:

- AI system inventory by jurisdiction, maintained as a living map rather than a point-in-time snapshot.
- Use-case risk classification that is internally coherent, so it can be mapped onto whatever external framework a jurisdiction adopts.
- Developer and deployer responsibility mapping across vendor relationships.
- Vendor contract review for compliance allocation and information-sharing obligations.
- Automated decision documentation that exists regardless of whether a specific law currently requires it.
- A bias testing evidence repository, maintained as an ongoing practice rather than as a response to any single statute.
- Consumer notice and appeal workflows that can be enabled on a per-jurisdiction basis as required.

- A designated owner for regulatory monitoring, accountable for tracking change and triggering updates.
- A quarterly governance refresh cycle calibrated to the regulatory calendar.

The common thread is that none of these capabilities is tied to a specific law. They are the operating infrastructure that enables an organization to respond to Colorado's reversal, Connecticut's expansion, and whatever happens next without having to rebuild from scratch each time. A static AI policy will not survive a dynamic regulatory environment. A living governance system will.

> ***"A static AI policy will not survive a dynamic regulatory environment. The capabilities that matter are the ones not tied to any single law, because any single law can be reversed."***

Board-Level Questions

For leadership teams, Colorado's reversal reframes the questions worth asking:

- Which AI systems do we operate in Colorado, and which influence consequential decisions under the new SB 26-189 framework?
- How much of our prior Colorado compliance work was tied to the first law, and how much transfers to the replacement?
- Can we determine which of our systems are in scope under each jurisdiction's current standard and adapt as those standards change?
- Do our vendor contracts clearly allocate developer and deployer obligations to survive a change in the underlying law?
- Do we have decision documentation and bias-testing evidence that exists independent of any single regulatory requirement?
- Can we apply the same governance process to Connecticut, Texas, Utah, California, New York City, and the EU, with jurisdiction-specific configuration rather than jurisdiction-specific programs?
- Who owns the regulatory change function for AI, accountable for tracking shifts and triggering updates before they become compliance gaps?

Closing

Colorado is not the exception. It is the preview.

The first comprehensive state AI law in the country was built, contested, delayed, litigated, and repealed within two years, before it ever took effect. That is the environment in which enterprise AI governance now operates: not a stable set of requirements to comply with, but a shifting, fragmenting, occasionally reversing landscape that no single compliance program can keep pace with.

The organizations that handle AI governance well will not be the ones with the longest policy documents or the most detailed single-jurisdiction compliance programs. They will be the ones with governance systems that can adapt as regulatory demands shift, in whatever direction they shift. Colorado just demonstrated that the direction is not always forward.

> ***"The organizations that handle AI governance well will not have the longest policy documents. They will have governance systems that adapt as expectations move, in whatever direction they move."***

Coming soon: *I will be writing on LogicGate Risk Cloud, which covers both AI governance and Third-Party Risk Management (TPRM), a topic I will be exploring further in future articles. Stay posted.*

AI Governance Series: ***The AI Governance Tightrope*** **|** ***The New Control Layer*** **|** ***The Operating Model for AI Governance*** **|** ***Three Frameworks*** **|** ***The Specification Owner*** **|** ***The Organizational Intelligence Layer*** **|** ***AI Governance Is Becoming an AI Security Problem*** **|** ***AI Governance Is Becoming a Multi-State Operating Model Problem***

About the Author

Greg Aldrich is a Global CIO and Strategic Advisor with 30+ years of experience helping boards, executive teams, and C-suite stakeholders navigate IT strategy, AI governance, and digital transformation. He serves as a fractional CIO and Senior Strategy Advisor at Blue Tree Technology Group and Field CIO/CTO and Senior Transition Architect at SDS Consulting, and currently serves as CIO at Andrew Wommack Ministries & Charis Bible College, where he chairs both the AI Committee and IT Steering Committee. He has advised organizations across financial services, gaming, healthcare, higher education, logistics, and nonprofit sectors.

Connect: linkedin.com/in/galdrich | Aldrich.ai